ESTE: Ensemble Spanning Tree-based Query Plan Enumeration in Database Systems

Yesdaulet Izenov[1, 2]
Asoke Datta[2], Brian Tsan[2], Abylay Amanbayev[2], and Florin Rusu[2]

[1]Nazarbayev University, Kazakhstan
[2]University of California Merced, USA

NUARC

INTRODUCTION

Query plan enumeration is a critical component of query optimizer. To find an optimal execution plan, conventional query optimizers employ either an exhaustive or heuristic plan enumeration strategy:

- Exhaustive strategies ensure optimality but are computationally expensive for queries with many joins.
- Heuristic strategies, while faster, often select suboptimal plans.

As data and workload demands continuously change and increase, query optimizers must possess flexibility and avoid suboptimal plans.

Shifting between these strategies can lead to increased optimization time and poor plan quality. Moreover, transitioning the optimizer to a different strategy entails significant development costs.

OUR APPROACH

In this work, we present a robust and flexible strategy based on spanning trees, bridging the gap between the two strategies. By redefining the problem of finding an optimal plan as finding low-cost spanning trees, we create a unified framework to utilize various spanning tree algorithms. Leveraging the polynomial time complexity of spanning tree algorithms, we introduce ESTE (Ensemble Spanning Tree Enumeration):

- ESTE employs Prim's and Kruskal's spanning tree algorithms together to enhance optimizer robustness by exploring diverse areas of the search space.
- Additional spanning tree algorithms also can be integrated. ESTE offers a cheaper alternative to maintain optimizer robustness without completely changing the optimizer.
- ESTE empirically shows better robustness for large queries with minimal overhead and consistent optimization time.

EXISTING APPROACHES

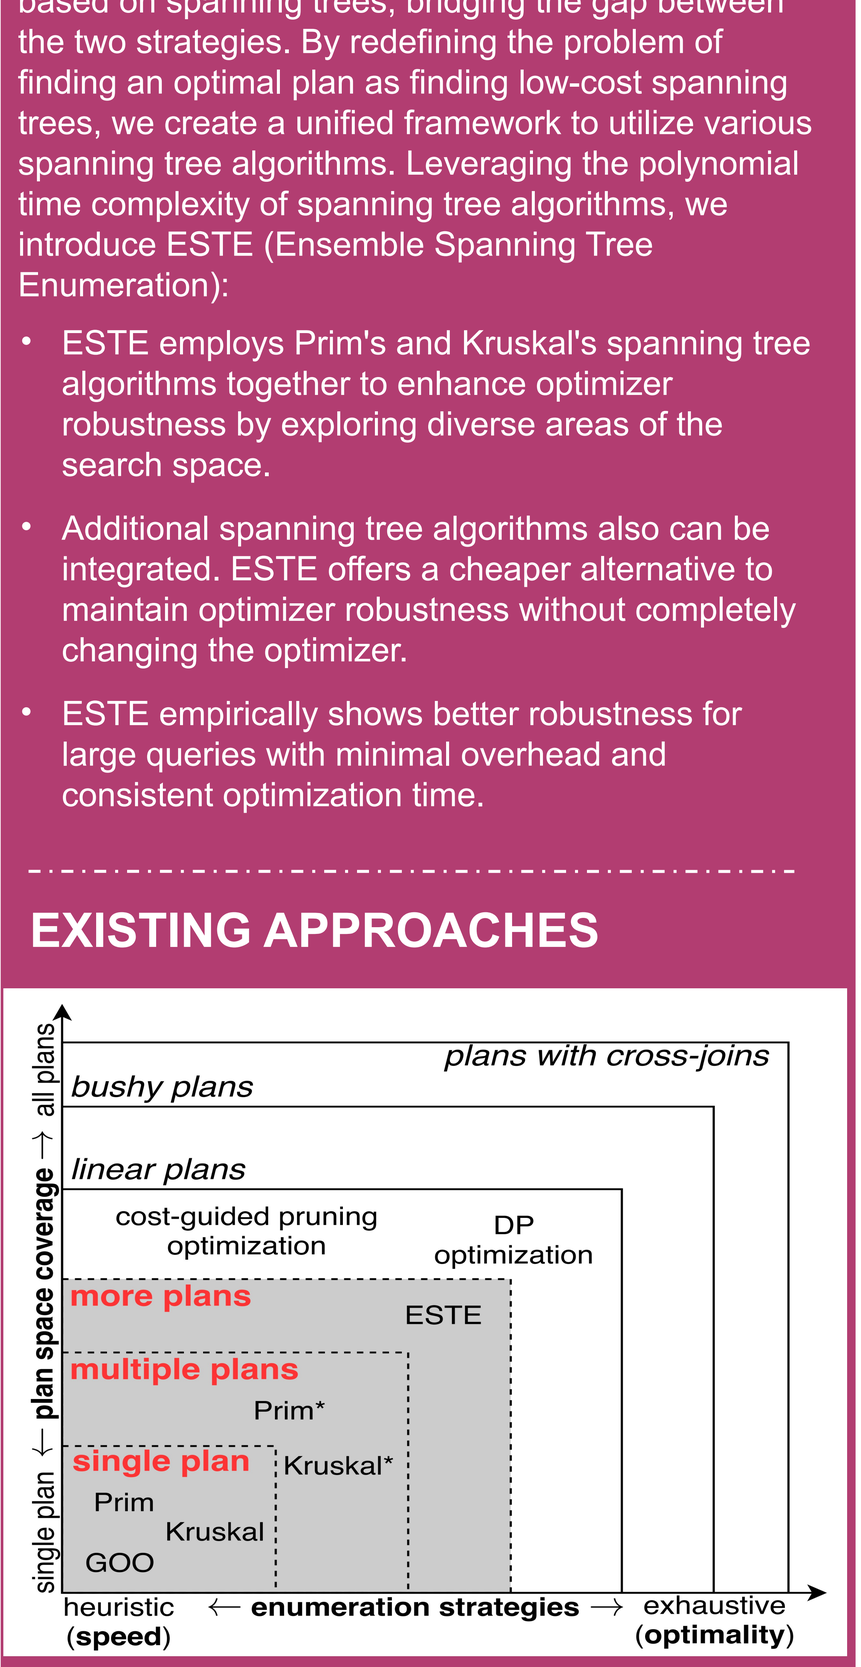

OUR TECHNICAL CONTRIBUTIONS

1. We introduce Ensemble Spanning Tree Enumeration (ESTE), a novel flexible approach that employs multiple spanning tree algorithms over the join graph. Thus, more plans from diverse areas of the search space are enumerated to increase the robustness of query optimizer.
2. As spanning trees are built, we adapt both Prim's and Kruskal's algorithms to account for changes in edge weights, multiple join and scan operators. Thus, the resulting trees are physical plans, with various binary tree shapes, rather than logical plans.
3. In our extensive evaluation, we first evaluate query plans produced by ESTE by comparing them against exhaustive enumeration and GOO which is recognized as one of top-performing heuristic methods.
4. We examine the performance behavior of ESTE in the presence of cardinality estimation errors. As expected, we observe a decline in efficiency across all enumeration methods, including exhaustive enumeration. The results exhibit higher overall performance of ESTE compared to the other considered enumeration methods, while also maintaining low optimization time.
5. We examine ESTE performance across different graph topologies - chain, cycle, start and clique - and compare ESTE to more well-known and recent heuristic methods.

A RESULTS: Query Plan Cost Comparison

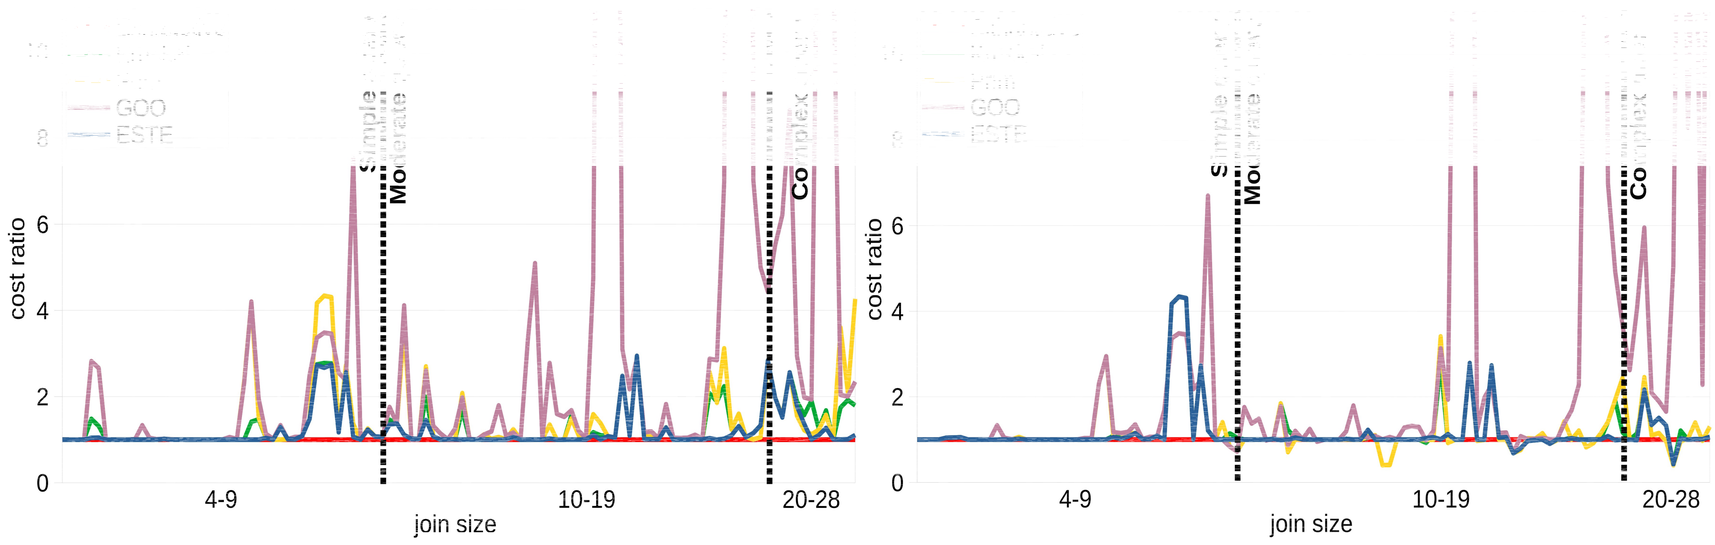

Cost ratio of selected and optimal query plans in Join Order Benchmark.
The costs are computed using true cardinalities.

B RESULTS: Query Execution Time Comparison

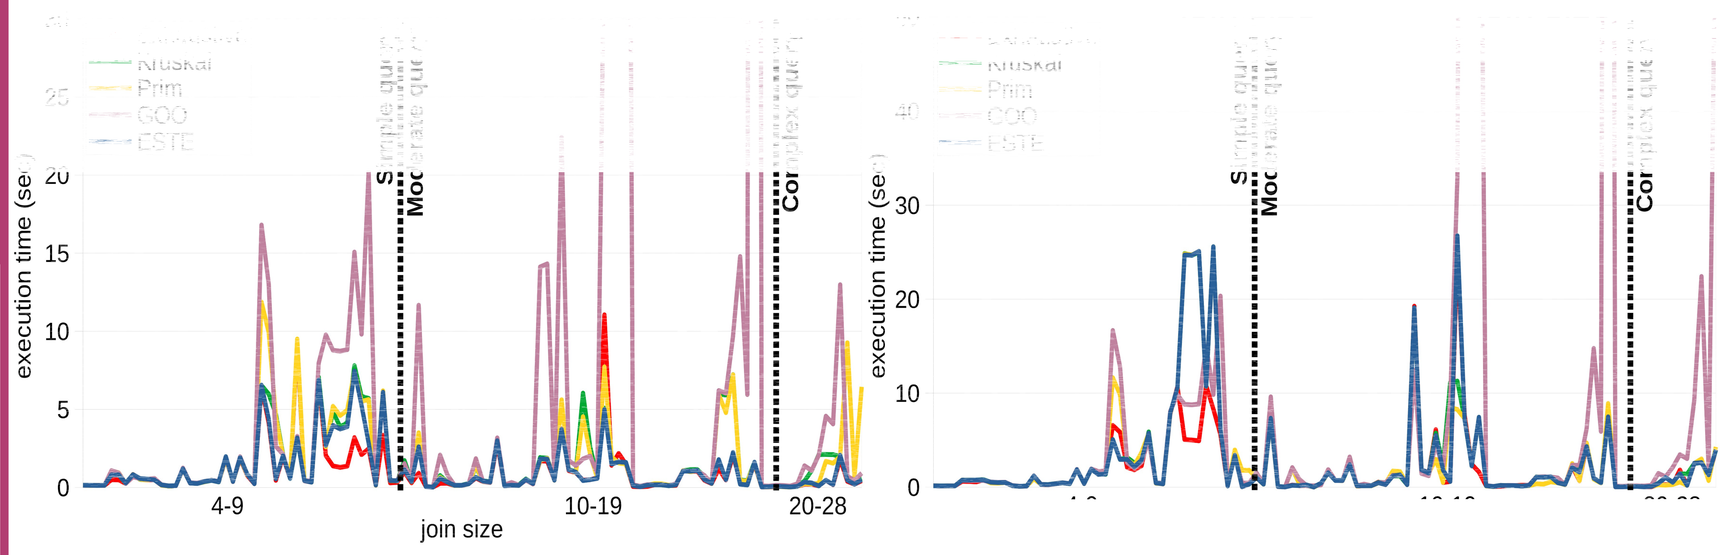

Execution time (sec) of query plans selected by enumeration algorithms.

Acknowledgment

U.S. National Science Foundation (NSF), Award number 2008815

References

Spanning Tree-based Query Plan Enumeration, 2024, in review
https://arxiv.org/abs/2403.04026

Sustainable Development Goals

9 INDUSTRY, INNOVATION AND INFRASTRUCTURE
6 CLEAN WATER AND SANITATION
11 SUSTAINABLE CITIES AND COMMUNITIES
15 LIFE ON LAND